Comments on AI Systems

One goal of course is to appreciate what is feasible with AI today—what is realistic and what is sci fi.

Many of systems presented involved an agent interacting with the world (sensors, actuators)

- humanoid robotics (see *Wired* article)
- contrast with traditional AI systems: microworlds, and/or focus on one aspect of cognition in isolation

Origins of AI

1943 **"A logical calculus of the ideas immanent in nervous activity" (McCulloch and Pitts)**

Boolean circuit model of brain

laying foundation for neural networks

1948 ***Cybernetics*** **(Norbert Wiener)**

control theorist wrote about biological and mechanical control systems

introduced public to notion of AI machines

Origins of AI (continued)

1950 Turing "Computing Machinery and Intelligence"

Can machines think? Can machines behave intelligently?

Operational test for intelligent behavior: imitation game

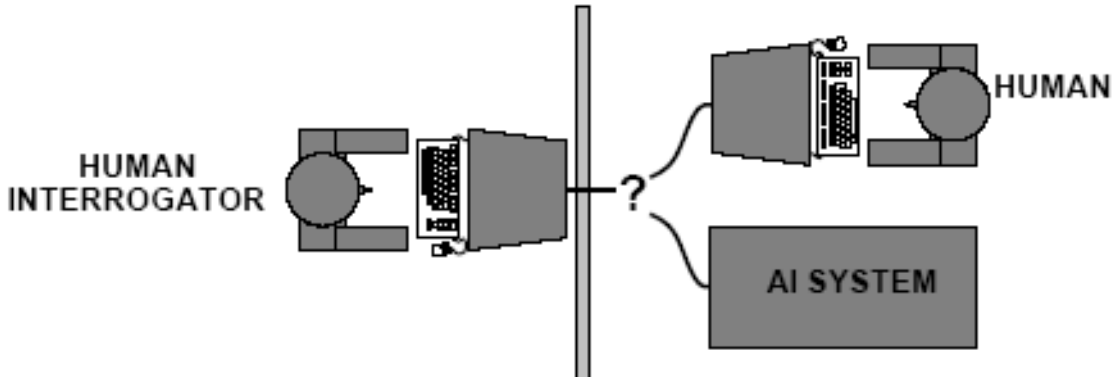

Predicted that by 2000, machine might have a 30% chance of fooling a lay person for 5 minutes

Suggested major components of AI: Knowledge representation, reasoning, language understanding, learning

1956 Dartmouth meeting

"Artificial Intelligence" adopted

Early Symbolic AI

1955 **Logic theorist (Newell, Shaw, & Simon)**

rederived theorems from *Principia Mathematica* (fundamental properties of numbers in terms of pure logic)

given: set of axioms of logic

derive implications (theorems) from axioms

heuristic search

1957 **General problem solver (Newell, Shaw, & Simon)**

domain independent

abstract rules of problem solving, e.g., “if you can, always try to bring your current state closer to the goal state”

did not scale

domain specific knowledge and heuristics were needed

Early Symbolic AI (continued)

1952-62 Samuels's checkers program

automatic tuning (learning) of heuristic evaluation function

1958 Invention of LISP (McCarthy)

1963 ANALOGY (Evans)

1964 Algebra word problems (Bobrow)

early natural language understanding

1965 ELIZA (Weizenbaum)

interactive program, carries on a dialog on any topic

Rogerian psychotherapist version very popular

http://www.manifestation.com/neurotoys/eliza.php3

http://www.pandorabots.com/pandora/talk?botid=f5d922d97e345aa1

Early Symbolic AI (continued)

1965 **resolution theorem proving for first-order logic (Robinson)**

complete: guaranteed to draw conclusion if it follows from knowledge base

limitation: scaling properties

1966 **semantic networks (Quillian)**

simple knowledge representation scheme

1966 **negative report on machine translation**

kills much work in NLP

1969 **Shakey (SRI)**

robot combining locomotion, perception, and problem solving

1971 **SHRDLU (Winograd)**

robot arm carries out instructions typed in English in blocks world

Early Symbolic AI (continued)

1972 **Development of Prolog (Coimerauer)**

theorem proving language

1979 **Nonmonotonic logic and truth maintenance systems (McDermott & Doyle; McCarthy)**

1980 **First AAAI (American Association for Artificial Intelligence) conference**

More Recent Symbolic AI

1969-79 Early development of knowledge-based systems

Realization that AI systems build on general-purpose search and inference mechanisms were not powerful enough to solve difficult problems.

Knowledge of specific domains is key to success.

DENDRAL: infers molecular structure from information provided by mass spectrometer

MYCIN: medical diagnosis (blood infections)

1980-88 Expert systems industry booms

R1: configures orders for new DEC computer systems

1988-93 Expert systems industry busts; "AI winter"

Early Subsymbolic AI

1949 ***The organization of behavior*** **(Hebb)**

proposed mechanism by which the brain adapts

neurobiological reality has since been verified

1962 **Perceptrons (Rosenblatt)**

A perceptron can learn to do any task it is programmed to do.

1969 **computational complexity analysis of perceptrons (Minsky and Papert)**

computational complexity analysis of perceptrons

origin of complexity analysis

1970-85 **neural networks abandoned**

1970s **computer vision (Marr)**

reconstructing 3D representation of world from camera images

Modern Subsymbolic AI

1982 **attractor networks (Hopfield)**

Neural network computing can be viewed as optimization

1986 **back propagation (Rumelhart, Hinton, & Williams)**

Neural network learning can be viewed as optimization

1988 ... **resurgence of probabilistic and statistical methods in AI**

genetic algorithms, Bayesian networks, support-vector machines

Integration of ideas from symbolic and subsymbolic AI (e.g., Bayesian networks, artificial life, intelligent agents)

Issues

Scaling

Symbolic versus subsymbolic AI

Human versus rational intelligence

human intelligence: fidelity to human performance

rational intelligence: ideal performance; makes best possible inferences, performs best possible actions

For example...

Sam is a construction worker.

Sam is a construction worker who drives a pick up truck and goes out for a beer with his buddies after work on Friday.

Bounded (limited) rationality

limitations on knowledge, time, computational power

“satisficing” (good enough) solutions